

AFRL-IF-RS-TR-2001-256
Final Technical Report
April 2002



ESTABLISHMENT OF CENTER OF EXCELLENCE IN MULTISOURCE INFORMATION FUSION

Calspan – UB Research Center, Incorporated (CUBRC)

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

20020610 034

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.


AFRL-IF-RS-TR-2001-256 has been reviewed and is approved for publication.

APPROVED:



DANIEL R. KUPIAK
Project Engineer

FOR THE DIRECTOR:



JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/IFEE, 26 Electronic Parkway, Rome, NY 13441-4514. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE APRIL 2002		3. REPORT TYPE AND DATES COVERED Final Dec 96 - Dec 00
4. TITLE AND SUBTITLE ESTABLISHMENT OF CENTER OF EXCELLENCE IN MULTISOURCE INFORMATION FUSION			5. FUNDING NUMBERS C - F30602-96-C-0336 PE - 062702F/062204F PR - 1013 TA - 10 WU - P1	
6. AUTHOR(S) James Llinas, Chris Bowman, Galina Rogova, and Henry Hexmoor				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) CALSPAN - UB Research Center, Incorporated (CUBRC) PO Box 400, 4455 Genesee Street Buffalo New York 14225			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFEE 26 Electronic Parkway Rome New York 13441-4514			10. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2001-256	
11. SUPPLEMENTARY NOTES Air Force Research Laboratory Project Engineer: Daniel R. Kupiak/IFEE/(315) 330-3206				
12a. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This Final Technical Report is a summary of the work conducted on the following research tasks: (1) Correlation Science, (2) Learning in Distributed Fusion Systems for Decision Making and (3) Information Sharing in Distributed Fusion Applications. The Correlation Science task has derived a new association scoring approach, JHSO, and compared its performance with four traditional association approaches. The JHSO yields the best association and estimation results for problems where an accurate report and track are being considered for association versus another less accurate track (e.g. from off-board). The other approaches are biased towards less accurate tracks since they do not consider the resultant state accuracy in their scoring for alternative associations. The Learning in Distributed Fusion Systems for Decision Making task addressed issues of cooperative group learning from examples in multi agent distributed system for decision making. An approach to both individual learning with decision integration and distributed learning utilizing the Dempster Shafer theory of evidence has been introduced. More research is needed to address fundamental issues of the problem of cooperative learning in problem solving systems such as utilization of a prior knowledge, incorporation of symbolic and numeric information, and convergence of the process. The Information Sharing in Distributed Fusion Applications task presented a model and testbed for situation assessment. Many parameters affecting information sharing and strategies for information sharing were considered but a few key parameters are still missing. A trade off of 15% detection rate for a 30% saving in communication was shown. Cost and other factors such as models of other agents need to be included in the testbed to make it more useful.				
14. SUBJECT TERMS Information Fusion, Information Technology, Data Fusion, Data Association, Reinforcement Learning, Dempster-Shafer Methods			15. NUMBER OF PAGES 32	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

Table of contents	i
Foreword	ii
Acknowledgement	ii
I. Research on Correlation Science	1
1.0 Definition of Scoring Metrics	1
1.1 Chi-Square (CHI) Score	2
1.2 Integral of the Tail of the Chi-Square Distribution (CHI-Tail) Score	2
1.3 Max A Posteriori (MAP) Score	2
1.4 Probabilistic Data Association (PDA) Score	4
1.5 Joint Association and State Optimization (JHSO) Score	4
2.0 Performance Evaluation Criteria	5
3.0 Performance Evaluation Results	5
4.0 References	10
II Learning in Distributed Fusion Systems for Decision-Making	11
1.0 The Model	11
2.0 Results	13
3.0 Conclusions	14
4.0 References	15
III Information-Sharing in Distributed Fusion Applications	17
1.0 Introduction	17
2.0 The SAM-ISS Testbed	18
3.0 Test Scenarios	19
4.0 Results and Discussion	21
5.0 Conclusion	22
6.0 References	23

LIST OF FIGURES

Figure 1: A Comparison of MAP, JHSO, and PDA Hypothesis Evaluation	3
Figure 2: Comparison of Alternative Gaussian-Based Association	4
Figure 3: % Misassociation Comparison with sigS=1	7
Figure 4: Average Performance Comparison with sigS=1	7
Figure 5: % Misassociation Comparison with sigS=1	8
Figure 6: % Misassociation Comparison with sigS=1	8
Figure 7: Average Performance Comparison with sigS=1	9
Figure 8: Average SD Comparison (object 3) with sigS=1	9
Figure 9: Distributed System Architecture	12

Foreword

This Final Technical report is a brief summary of the work conducted on each of three research tasks at the Center for Multisource Information Fusion (CMIF) at the University at Buffalo, Buffalo, New York¹ over the period 1996-1999. Under co-sponsorship of the Air Force Research Laboratory's (AFRL) Information (IF) and Sensors (SN) Directorates, CMIF conducted research on three different topics in this period:

- Research on Correlation Science
- Learning in Distributed Fusion Systems for Decision-Making
- Information-Sharing in Distributed Fusion Applications

This report is a compilation of some of the key points in terms of purpose and motivation, as well as results, associated with each of these tasks. Individual technical reports resulting from each task have been previously sent to AFRL.

Acknowledgement

CMIF gratefully acknowledges the support of the Air Force Research Laboratory for the conduct of this research.

¹ The "University at Buffalo" is the current name for what has often and heretofore been known as the State University of New York at Buffalo.

I. Research on Correlation Science

Principal Investigators: Dr. James Llinas, Research Professor, UB; Dr. Chris Bowman, Consultant;
Selcuk Savas, Graduate Student

The emphasis of this project was selected to be the function of data fusion that has the greatest impact on operational system performance. The two key factors for such a selection is the importance of the function and its maturity. The least mature of all the functions of data fusion is data association, which has three sub-functions (i.e., hypothesis generation, evaluation, and selection). Hypothesis generation is a more *ad hoc* function that reduces the size of the rest of the data association problem. Hypothesis selection is the most mature of the three functions due to the extensive search algorithm developments performed in the field of operations research. This effort thus focused on the hypothesis evaluation function of data association since it is the least mature and is the basis upon which hypothesis selection and the rest of the fusion problem is solved. Within that function emphasis was placed on the approach to and nature of the "scoring metrics" used within the hypothesis evaluation processing.

1.0 Definition of Scoring Metrics

The performance of the following five scoring metrics is evaluated in a selected set of test cases:

1. chi-square (CHI) statistic
2. integral of the tail of the chi-square distribution (CHI-Tail)
3. max a posteriori (MAP)
4. probabilistic data association (PDA)
5. joint hypothesis and state optimization (JHSO)

This section refines these five metrics for three cases of 1-dimensional report-to-track problems as follows:

1. case 1: the association of a single report with each of two tracks having different variances (i.e., cases 1a and 1b)
2. case 2: a report on a new object (e.g., a pop-up target) not associated with either of two existing tracks
3. case 3: 2 reports associated with 2 tracks having different variances

1.1 Chi-Square(CHI) Score

The CHI score (i.e., also known as the Mahalanobis score) is the exponent of the Gaussian for the difference between the report, S, and the track, T. Namely,

$$\text{CHI score} = \mathbf{I}^T \mathbf{V}^{-1} \mathbf{I} = (\mathbf{S} - \mathbf{T})^2 / [\mathbf{R} + \mathbf{P}]$$

where R is the variance of S and P is the variance of T. The association hypothesis selected is the one that minimizes this score. For the multi-report case, sum these scores is minimized. For case 2 a 3-sigma limit for association is used, so the chi score for non-association is 9.0.

1.2 Integral Of The Tail Of The Chi-Square Distribution (CHI-Tail) Score

The integral of the tail of the chi-square distribution (CHI-Tail) is used in statistical hypothesis testing approaches for deterministic data association. To compute the of integral of tail of chi square, for one dimension first compute the integral of the N(0,1) Gaussian from negative infinity to the square root of the CHI score, x, (i.e., $N(\sqrt{x})$) and insert into the equation below:

$$\begin{aligned} \text{CHI-Tail} &= 1 - \text{int of tail of chi}(x) \\ &= 1 - .8 [\sqrt{2\pi}] [N(\sqrt{x}) - .5] \end{aligned}$$

For example, the integral of tail score (where the current chi score = .455) is $\text{CHI-Tail}(.455) = 1 - .8 [2.5] [N(.675) - .5] = 1 - 2 [.75 - .5] = .5$. Choose the association that has the largest CHI-Tail score. For case 2 the equivalent of a 3-sigma limit for association is used, so the integral of the tail score for non-association is 0.0026. For the single report case 1 experiments herein the association results will be the same as with CHI, since where two CHI scores are equal two CHI-Tail scores will also be equal.

1.3 Max A Posteriori (MAP) Score

There are various max a posteriori (MAP) scoring criteria for data association and state estimation [4]. The most common of these is the one defined on top of Figure 1. This the usual MAP deterministic data association criterion used to select the "best" hypothesis, which is then used to generate the MAP estimate of the system state. The second defined in the Figure is probabilistic data association (PDA), which updates the track state confidence for each report based upon its relative association confidence score. The third criterion in the figure is the joint hypothesis and state optimization (JHSO) score.

The MAP kinematic scoring for a new incoming sensor report, S, to an existing track, T, assumes a Gaussian distribution with a central track covariance P that models the error in the track location. For the single report case 1a and 1b and for the two-report case the non-association hypotheses are assumed to be near zero probability. For case 2 we test an input report that is a pop-up (i.e., not the same object as any of the tracks). For this case the scene hypothesis scores are each the product of the individual hypothesis scores for how all the given batch of reports and tracks are associated. Namely,

- DETERMINISTIC DATA ASSOCIATION THEN TARGET —

$$\underset{H}{\text{MAX}} P(H | \text{REPORTS}) = \underset{H}{\text{MAX}} [P(\text{REPORTS} | H) P(H) \text{ THEN } \underset{\theta}{\text{MAX}} P(\theta | \hat{H})]$$
- TARGET STATE ESTIMATION WITH PROBABILISTIC DATA

$$\underset{\theta}{\text{MAX}} P(\theta | \text{REPORTS}) = \underset{\theta}{\text{MAX}} \left[\sum_H P(\text{REPORTS} | H, \theta) P(H | \theta) \right] P(\theta)$$
- JOINT ASSOCIATION DECISION AND TARGET STATE

$$\underset{H, \theta}{\text{MAX}} P(H, \theta | \text{REPORTS}) = \underset{H}{\text{MAX}} [\underset{\theta}{\text{MAX}} P(\theta | \text{REPORTS}, H)] P(H | \text{REPORTS})$$

Figure 1: A Comparison of MAP, JHSO, and PDA Hypothesis Evaluation Criteria

1. Association Hypotheses:

$$P(H|S,T) = P(S|T,H) P(H) = \{|V|^{-1/2}\} \exp[-1/2 \{I^T V^{-1} I\}] [1-P_{FA}(S)] [1-P_{FA}(T)] P_D(S) P_D(T)$$

2. Pop-up and Track Propagation Hypotheses:

$$P(H|S,T) = P(S|T,H)P(H) = \{E(|P+R|^{-1/2})\} \exp[-1/2\{.455\}][1-P_{FA}(S)][1-P_D(T)]P_D(S)[1-P_{FA}(T)] [1-P_D(S)] P_D(T)$$

where for these evaluations,

- S is the sensor report with variance R,
- T is the track with variance P,
- H is the hypothesis that the report and track are associated,
- |V| is the determinant of the innovations covariance, $V = P + R$,
- I is the innovations vector, $I = S - T$
- $P_D(S)$ is the probability of detection for this object reported by the sensor = .5 here.
- $P_{FA}(S)$ is the probability of false alarm (FA) of the sensor for this type of report = .1
- $P_D(T)$ is the probability of detection of this object in the track file = .7
- $P_{FA}(T)$ is the probability that this track is a false alarm = .01
- the $E(|P+R|^{-1/2})$ for one report (i.e., S) and two tracks (i.e., T1 and T2) is the following average

$$\{|V(S-T1)|^{-1/2} + |V(S-T2)|^{-1/2}\}/2$$
- False Alarm and Track Drop hypotheses are assumed to have a much lower confidence so are not used here.

The normalized comparison of these first three scores on a N(0,1) distribution, given in Figure 2, shows the differences in the scores as the report-to-track separation increases. This difference emphasizes the importance of selecting the correct kinematics-scoring scheme to be rigorously comparable with the ID and a priori statistics.

Report-Track Error	Gaussian MAP	χ^2 Integral of Tail	CHI (Mahalanobis)
0 σ	1.0	1.0	0.
.1 σ	.995	.92	.01
.32 σ	.95	.75	.1
.4 σ	.92	.7	.16
.675 σ	.796	.5	.455
1 σ	.6	.32	1.
1.15 σ	.5	.25	1.32
1.6 σ	.275	.12	2.6
2 σ	.13	.04	4.

Figure 2: Comparison of Alternative Gaussian-Based Association Scoring Techniques

1.4 Probabilistic Data Association (PDA) Score

The PDA score maximizes the probability of the object state given the reports. Thus, the PDA updated state estimates are as follows:

- For case 1: $[P(\text{association}) = 1]$
T1 estimate = [Kalman update of T1] [MAP association score S to 1] + draw for T1 [1 - MAP association score S to 1]
T2 estimate = [Kalman update of T2] [MAP association score S to 2] + draw for T2 [1 - MAP association score S to 2]
- For case 2: $[P(\text{false alarm \& track})= 0]$
S estimate = [Kalman update of T1 with S] [MAP score S to 1] + [Kalman update of T2 with S] [MAP score S to 2] + draw for S [MAP score that S not associated]
T1 estimate = [Kalman update of T1 with S] [MAP score S to 1] + draw for T1 [1 - MAP score S to 1]
T2 estimate = [Kalman update of T2 with S] [MAP score S to 2] + draw for T2 [1 - MAP score S to 2]
- For case 3: $[P(\text{association} = 1)]$
T1 estimate = [Kalman update S1 to T1] [MAP score of S1 to T1] + [Kalman update S2 to T1] [MAP score of S2 to T1]
T2 estimate = [Kalman update S1 to T2] [MAP score of S1 to T2] + [Kalman update S2 to T2] [MAP score of S2 to T2]

1.5 Joint Association and State Optimization (JHSO) Score

Instead of maximizing over all the hypotheses, as in the MAP, or over all state estimates, as in the PDA, this score maximizes over both simultaneously. Namely, the JHSO association hypothesis kinematic scoring for a new incoming sensor report, S, to an existing track, T, each with a Gaussian distribution having a report and track covariance of R and P(k-1), respectively is computed for association hypotheses as follows:

$$P(\theta|S,T,H) P(H|S,T) = \{|P(k)|^{-1/2}\} \{|V|^{-1/2}\} \exp[-1/2 \{I^T V^{-1} I\}] [1-P_{FA}(S)][1-P_{FA}(T)]P_D(S)P_D(T)$$

where

- $\bar{\theta}$ is the MAP estimate of the track kinematics state updated with the associated report, $\bar{\theta} = T(k)$
- $T(k) = T(k-1) + K [S(k) - T(k-1)]$
- $|P(k)|$ is the determinant of the updated track kinematics state covariance,
 $P(k) = [I - K] P(k-1)$
- K is the track state update Kalman gain,
 $K = P(k-1) [P(k-1) + R]^{-1}$

2.0 Performance Evaluation Criteria

The performance evaluation criteria are the percentage of mis-associations for the four that make such a decision (i.e., all but PDA). For all approaches the mean and variance of the error in the Kalman estimate is computed. The combined average performance metric used is the root sum squared (RSS) of the means in the errors divided by the standard deviations in Kalman estimate given the correct association for each object. Namely,

$$\text{Avg. Perf.} = \{ \sum^j (\text{mean error}(j)^2 / P(1)(j)) \}^{1/2}$$

where the sum is over the true objects, j , and $P(1) = (I - K) P(0)$ is the updated covariance of the object states where K is the optimal Kalman gain.

3.0 Performance Evaluation Results

The performance results show that no single association approach is the "golden score". For cases where all the covariances are equal, all the association decisions are the same. Thus we focus on differing covariance cases and to keep the representations simple a single kinematics dimension is used following all the basic Kalman filter assumptions. Case 1 tests the association of a single report, S , with one of two tracks, $T1$ and $T2$. S is simulated to have arisen from the same object as $T1$ by performing a Gaussian draw for each at the same mean (i.e., $\mu_S = \mu_{T1}$). This mean is varied from -25 to 45; all draws for S and $T1$ have a standard deviation of 1. Track 2 is drawn from a Gaussian with mean and standard deviation of 10. Each trial at a given set of means is run thru 1500 draws to compute the average mis-association, mean error, and standard deviation of the mean error for each object.

The percentage of mis-associations for each of the first four association selection approaches (i.e., deterministic association scoring schemes) is shown in Figure 3. CHI and CHI-Tail have the same association decisions even though the shape of their scores with report-to-track differences varies. This is due to the fact that with only one report and 2 tracks the point at which the CHI scores are equal is the same point in all cases as where the CHI-Tail scores are equal (i.e., one is a function of the other). This is not the case in general. For this case 1, JHSO performs slightly better than MAP and both significantly better than the CHI scores. This improved association performance also yields better tracking performance as shown in Figure 4. In fact both JHSO and MAP have better tracking performance than PDA even though PDA bases its association confidence update weight on MAP. For comparison the results for the same set of

trials, but with the standard deviation of the draw for T2 equal 4, showed an even larger performance improvement of MAP and JHSO. For another comparison where the accuracy of the associated track is larger than that for the non-associated track, we made the mean of the report equal to the mean of track T2 (i.e., $\mu S = \mu T2$) and set $\mu T1 = 2$. The RSS tracking error performance is plotted in Figure 5. In this case the CHI scores perform better than both MAP and JHSO both for association and tracking performance, since the CHI scores favor the larger covariance track. As before, PDA falls in between.

One of the major benefits of the MAP and JHSO scores is that they are rigorously compatible with scores incorporating other information (e.g., probability of detection and false alarm, ID attributes, etc.). To provide a simple comparison of this performance we insert a report on an object that is not either that for T1 or T2. We insert this report with a standard deviation of 1, at various means from -12 to 22 and leave $\mu T1 = 2$ with $\sigma T1 = 1$ and $\mu T2 = 10$ with $\sigma T2 = 10$. This case 2 tests the ability of the 5 approaches to distinguish a new pop-up object. The mis-association and RSS tracking error results are shown in Figures 6 and 7, respectively. The CHI scores show different association decisions since there is now three association choices. The CHI-Tail performs slightly better than CHI and both are significantly worse than JHSO and MAP. MAP has the smallest RSS tracking error, however JHSO has the smallest average standard deviation in the tracking error for all mean trials and all objects. For example, this average standard deviation in the error for the new object (3) is shown in Figure 8 and at $\mu S = -12$ for object 2 it is 8 vs >10 for all the others and for object 1 it is .4 vs 1 for all the others.

The performance for a case 3 with two reports and two tracks and no allowed pop-ups or propagated tracks was run next. The results for 1500 draws with the means of report and track 1 at 2 with sigma of 1 and means of report and track 2 at 10 with sigma of 10 showed MAP performing best with 3% mis-associations and an RSS accuracy of .03. JHSO and CHI came in second with 33% and .15. CHI-tail in third with 40% and .25. PDA performed worst with an RSS accuracy of 1.4. For these single dimensional case 3 results the JHSO yields the same results as the CHI since the JHSO state terms cancel to within a constant with the MAP terms. This in general is not the case.

In summary, this research task has derived a new association scoring approach, JHSO, and compared its performance with four traditional association approaches. The JHSO yields the best association and estimation results for problems where an accurate report and track (e.g., from on-board sensors) are being considered for association vs another less accurate track (e.g., from off-board). The other approaches are biased towards the less accurate track since they do not consider the resultant state accuracy in their scoring for the alternative associations.




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8

4.0 References

1. Bowman, "Affordable Information Fusion Via an Open, Layered, Paradigm-Based Architecture", Proceedings of 9th National Symposium on Sensor Fusion, Monterey, CA, March '96.
2. Bowman and C.L. Morefield, "Multisensor Fusion of Target Attributes and Kinematic Reports", 3rd ONR/MIT Conference on C31, June 1980.
3. Bowman "The Data Fusion Tree Paradigm and It's Dual" Proceedings of 7th National Symposium on Sensor Fusion, invited paper, Sandia Labs, NM, March '94
4. Llinas, B. Neuenfeldt, L. McConnell, D. Bohny, C. Bowman, D. Hall, et al, "Studies and Analyses Within Project Correlation: An In-Depth Assessment of Correlation Problems and Solution Techniques", Proceedings of 9th NSSF, Monterey, CA, March '96.

II. Learning in Distributed Fusion Systems for Decision-Making

Principal Investigators: Dr. Galina Rogova, Adjunct Professor, UB; Raghu Menon, Graduate Student

The research of this task addresses issues of cooperative group learning from examples in a multi-agent distributed system for decision-making. An approach to both individual learning with decision integration and distributed learning utilizing the Dempster-Shafer theory of evidence is introduced. The implementation of this approach in different machine learning paradigms corresponding to different types of agents' interaction with their environment is also presented. The case study considered in the research has demonstrated significant improvements not only in overall performance of the system, but in the case of distributed learning, and also in recognition accuracy of individual agents. The developed method is especially important in real-life applications when construction of a large training data set is very expensive or impractical. The introduced adaptive fusion method for learning in distributed systems is not problem specific and can successfully be used for both military and non-military applications for target recognition, situation assessment, image retrieval, pattern recognition, medical imaging, etc.

The objective of this research is to develop an approach to improved decision making ability of a multi-agent distributed system as the result of learning. The interest in the distributed problem-solving system stems from the fact that many real-world problems can be better solved using a set of cooperative agents rather than a single agent [1]. A general multi-agent system consists of a group of distributed intelligent agents that coordinate their knowledge, goals, skills and plans so that they can make decisions, take actions and solve problems (see, e.g., [1-4]). Interaction between agents may occur at a number of levels: at the knowledge level, at the planning level, or at the action level (knowledge, possibility, and choice) [5]. Agents in distributed systems may have different areas of expertise, *a priori* knowledge, and problem-solving abilities. They may be able to observe only certain characteristics of the environment, or they may observe different parts of the environment (spatially or temporally), or some of them may be able and some of them may not be able to interact with environment. Since no one agent has complete information about environment, they have to cooperate to achieve their goals. In the research addressed on this task, efforts were concentrated on the knowledge level interaction between agents and address the problem of learning in a distributed system that is designed to make decisions about the states of environment. The learning goal here is to adjust the system's decision-making process in order to improve its performance in future situations. In distributed systems, this goal can be achieved as a result of interaction between the system and environment and cooperation between the agents that may exchange their *a priori* knowledge, observations and local decisions.

1.0 The Model

This section addresses the issues of group learning in a multi-agent distributed system when agents learn individually and do not communicate during the learning process with each other. We consider here a two-layer distributed system where the lower level agents have different expertise and either can observe independent features of the environment, or use different processes to make local decisions about the states of environment. We assume here that

the set of hypotheses about environment is given to each agent *a priori*. This situation may appear when we have, for example, agents represented by different types of sensors, or agents extracting different features of environment. The lower level agents in our system are modeled by the artificial neural networks that are very effective when used for "low-level" matching and recognition. They are designed to learn from examples and are very robust when deal with uncertain and incomplete input information (both numeric and symbolic). We will consider the different types of interaction between the lower level agents and environment that correspond to various learning paradigms. In supervised learning, the environment can act as a knowledgeable teacher and provide explicitly the desired decision as a feedback which then is used to modify the decision- making scheme of the classifier. In reinforcement leaning, the environment evaluates the decision and provides only binary reinforcement signal about correctness of the decision. In unsupervised learning, the system does not obtain any feedback and changes the decision function based on internal structure of the training set. At the problem-solving stage, each agent uses its expertise to acquire information about the state of environment, makes local decisions about global hypotheses and transmits the decisions with their confidence levels to the fusion center where the agents decisions are combined in order to produce the final decision. The fusion center does not interact with environment and employs only the information produced by each individual agent. The fusion process combines decisions made by the agents by taking advantage of the correct decisions produced by each agent and dealing with conflicting decisions. The architecture of the system considered in our research is presented in Figure 9.

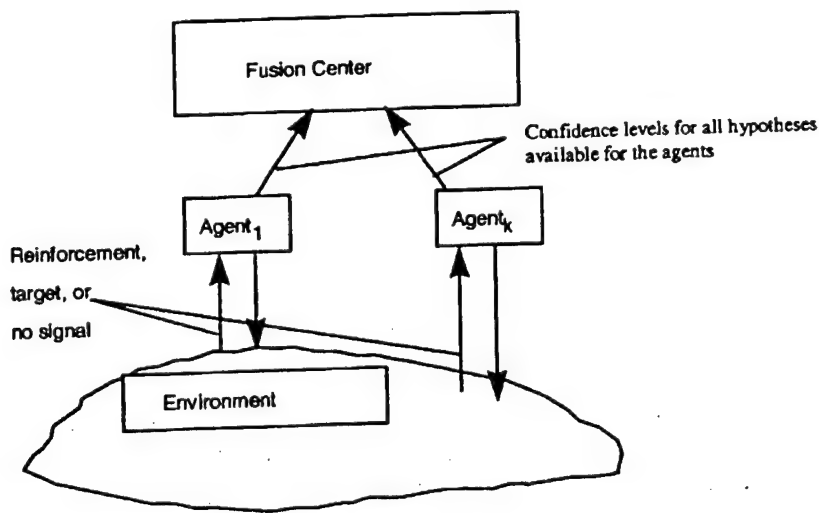


Figure 9 Distributed System Architecture.

The described process of individual learning by agents with subsequent decision integration in the fusion center has several advantages as compared to the process of distributed information gathering and individual learning when the agents make observations, transmit the observations to the fusion center for the learning based on composite observations. One of the advantages is that transmission of the confidence levels reduces the amount of information transmitted as compared to transmission of observations. Another advantage of employing a

decision fusion process is that it improves the overall problem solving quality of the system as compared with problem solving quality of an even the best individual agent. This process can also achieve better recognition accuracy with less number of training patterns that is especially useful when the number of training patterns is insufficient and the process of obtaining additional training patterns is costly or impractical or training time is limited. Several techniques have been used for decision fusion: majority voting [11,12], Bayesian decision theory [11], Neural Networks [13], the Dempster-Shafer theory of evidence [11,14,15]. The latter seems to be best suited for decision combination because it does not require many training patterns, and provides conflict resolution capability and means to include subjective judgment. It also allows us to represent ignorance and uncertainty.

The idea behind the process introduced in this report is borrowed from the well-known and widely used in pattern recognition "minimum distance classifier". In this classifier, a classification decision is based on the minimum distance in the feature space from a pattern in question to a class reference vector. In our evidential fusion system, a classification decision is based on the belief in any classification hypothesis that represents a combination of individual beliefs estimated for the neural networks modeling the agents. Individual beliefs are computed within the framework of the Dempster-Shafer theory of evidence as a function of a distance between pattern representation and output characteristics of the neural networks (output vectors or weight vectors depending on the type of learning). Belief combination is carried out with the Dempster rule of combination. Therefore, we can consider our evidential classifier as a combination of evidential distance classifiers (EDC) built upon NNs considered in the combination.

2.0 Results

The goal of the experiments described in this subsection is to show that the performance of the introduced evidential classifier that integrates decision of agents trained individually is better than:

1. Performance of the best individual agent.
2. Performance of the system where the fusion center modeled as a neural network is trained on combined observations transmitted by the lower level agents.

We have conducted three types of experiments corresponding to different kinds of interaction with environment by lower level agents: feedback as a reference signal (supervised learning), feedback as a reinforcement signal (reinforcement learning), and no feedback at all (unsupervised learning). The agents modeled as supervised, unsupervised, and reinforcement neural networks, respectively, are trained individually and then their decisions are combined in the fusion center. We also trained the fusion center modeled by a neural network with a long feature vectors constructed as a concatenation of the available feature vectors as input. The latter corresponds to the type of learning where the agents make observations and transfer them to the fusion center which is then trained on a feature vector containing all the observations (individual learning with distributed information gathering).

The results of learning on combined observations turned out not better than the results of the best individual agent. This phenomenon can be explained by the fact that we do not have the number of training patterns in our database that are sufficient for training with a higher dimensional input vector. We may hope that this result can be improved by adding new patterns for training the neural network, although adding new patterns to the training set may be very difficult, costly, and even impractical in many real-life situations. Utilization of the decision fusion method introduced in this research gives us a significant improvement in accuracy without incorporation of additional training patterns. The data shows the superior performance of group learning in distributed system with the developed decision fusion method as compared with performance of the individual learning in distributed system based on combined observations and performance of the best individual agent.

The results of additional experiments showed that the research on the distributed reinforcement learning method designed on the project yield significant improvement in the system performance. Overall system performance improves on average by 68% as compared with simple decision fusion with knowledge integration in the case when the lower level agents do not have access to environment. The distributed reinforcement learning technique may also improve the performance of the individual agents trained as reinforcement neural networks individually. For the "texture" and "color histogram" agents of these experiments, cooperative learning improves performance by an average accuracy over individual agents by 50%.

3.0 Conclusions

The research described in this report has addressed issues of cooperative group learning from examples in multi-agent distributed system for decision-making. An approach to both individual learning with decision integration and distributed learning utilizing the Dempster-Shafer theory of evidence has been introduced. The implementation of this approach in different machine learning paradigms corresponding to different types of agent interaction with environment has also been presented. The case study considered in the research has demonstrated significant improvements not only in overall performance of the system, but in the case of distributed learning, also in recognition accuracy of individual agents. The developed adaptive fusion method for learning in distributed systems is not problem specific and can successfully be used for both military and non-military applications for target recognition, situation assessment, image retrieval, pattern recognition, medical imaging, etc. The developed method is especially important in real-life applications when construction of a large training set is very expensive or impractical. However, more research is needed in order to address fundamental issues of the problem of cooperative learning in problem solving systems such as utilization of *a priori* knowledge, incorporation of symbolic and numeric information, convergence of the process, etc. A very important problem to be addressed in the future is the definition of measures of effectiveness and designing optimal learning strategies and method for knowledge representation and exchange.

4.0 References

1. Weis, G. Distributed Reinforcement Learning, *Robotics and Autonomous Systems*, 15, 1995, 135-142.
2. Sian, S.S. Adaptation based on cooperative learning in multi-agent systems, *Decentralized AI-2*, Y. Demazeu & J-P Muller (Eds.), Elsevier Science, 1991, 257-271.
3. Decer, K.S., Distributed Problem Solving Techniques: A Survey, *IEEE Trans. on System, Man, and Cybernetics*, Vol. 17, No. 5, 1987, pp.729-740.
4. Shaw, M.J., Harraw, B., Herman, S, Distributed artificial intelligence for multi-agent problem solving and group learning, in *Proc. of the twenty-Fourth Annual Hawaii International Conference on System Science*, Vol. 4, 1991, 13-26.
5. Muller, J-P and Demazeu, Y, *Decentralized AI*, Proceedings of European Workshop on Modeling an Autonomous Agents in Multi-Agent World, Cambridge, 1989.
6. Brazdil, P., *et al*, Learning in Distributed Systems and Multi-Agent Environments, in *Proc. of the European Working Session on Learning*, Springer-Verlag, pp. 412-423.
7. Shaw, M.J., Fox, M. Distributed Artificial intelligence for group decision support, *Decision Support Systems*, 9, 1993, 349-367
8. Tan, M. Multi-agent Reinforcement Learning: Independent versus Cooperative Agents, In *Proc. of the 10 th International Conference On Machine Learning*, pp. 330-337, 1993.
9. Towell, G., *et al* Refinement of Approximate Domain Theories by Knowledge-based Neural Networks, in *Proc. of the 8 th National Conference on Artificial Intelligence*, pp.723-728, 1990.
10. Shaw, M.J., Winston, A.B., Learning and Adaptation in Distributed Artificial Intelligence, *Distributed Artificial Intelligence*, V2, L. Gasser and M. Huhns (eds.), Pitman, London, 1989, 413-430.
11. Xu, L. Krzyak, A., Suen, C. Y., "Methods of combining multiple classifiers and their applications to handwriting recognition", *IEEE Trans. on System, Man, and Cybernetics*, Vol. 22, pp. 418-435, 1992.
12. L. Lam and Suen, C.Y., Application of majority voting to pattern recognition: an analysis of its behavior and performance, *Trans. on System, Man, and Cybernetics*, Vol. 22, pp. 418-435, 1992.
13. R. Maclin, R., Shavlik, J., Combining the Prediction of Multiple Classifiers: Using Competitive Learning to Initialize Neural Networks, in *Proc. of 14 th Int. Conf. on Artificial Intelligence*, 1995

14. Rogova, G. Combining the results of several neural network classifiers, *Neural Networks*, Vol. 7, pp. 777-781, May 1994.
15. Rogova, G. The Dempster-Shafer theory of evidence for classification of medical images. Proc. of the 6th *IEEE Dual-Use Technology and Application. Conf.*, 1996.
16. Shafer, G., *A Mathematical Theory of Evidence*, Princeton, MIT Press, 1976.
17. Barnett, J.A., Computational methods for mathematical theory of evidence, in Proc. of the 7th *Int. Conf. on Artificial Intelligence*, pp. 868-875, 1981.
18. Yoh-Han Pao, *Adaptive pattern recognition and neural networks*, Addison-Wesley, 1989
19. Menon, R. and Acharya, R., Content based image query from image databases using spatio-temporal transforms and fractal analysis methods, *IEEE International Conference on Image Processing*, September 1996.
20. Ballard, D., *An Introduction to Natural Computation*, The MIT Press, 1997.
21. Lynne, K. Competitive Reinforcement Learning, in Proc. of the 5th International Conference on Machine Learning, 1988, pp.188-199.
22. Vasilakos, A. and N. Loukas, N. "ANASA - A stochastic reinforcement algorithm for real-valued neural computation," *IEEE Transactions on Neural Networks*, Vol. 7, No. 4, 1996, pp. 830-842.
23. T. Kohonen, J. Hynninen, J. Kangas, J. Laaksonen, K. Torkkola, "LVQ_PAK: The Learning Vector Quantization Package Version 3.0", April 1995, Helsinki University of Technology, Finland.

III. Information-Sharing in Distributed Fusion Applications

Principal Investigators: Dr. Henry Hexmoor, Adjunct Professor, UB; Dr. James Llinas, Research Professor, UB; Elizabeth Cuddihy, Graduate Student

When situation assessment is performed individually by cooperative, friendly agents working toward a commonly held goal, optimal strategies for sharing information are needed. The best strategies are often domain dependant with respect to constraints imposed by the domain itself. Appropriate strategies for information sharing can be selected by looking at the overall performance of a strategy based on a set of predefined metrics measuring qualities such as information consistency across the agents, information completeness among the agents, and timeliness of information availability to each agent. This task conducted research focused on exploring the development of specific strategies for information-sharing among distributed fusion nodes, using a case-study approach. We have developed a prototype system, SAM-ISS, within a Theater Missile Defense scenario that we have used to investigate the utility of a set of metrics for determining the quality of three different information sharing strategies (ISS).

1.0 Introduction

The process of situation assessment is a common human activity. Within the cycle of perceiving one's surroundings, cognitively processing the input with other relevant information, followed by producing an appropriate action, situation assessment involves the processes of perception followed by understanding. These are difficult cognitive processes to model computationally yet these are essential components of an autonomous agent, be it human, an animal, or artificial.

Smith and Sage [9] define situation assessment as "the process of detecting and defining an opportunity or problem." Although this definition alludes to the necessity of cognitive processing to realize the situation, the definition does not take into account the act of performing data aggregation and fusion over input from an array of different modalities. For this reason, we will define situation assessment as **the mental process of aggregating sensory, non-sensory, and a priori input to construct a mental representation of a situation**. A situation is any meaningful abstraction of personal, group, and environmental information relevant to the agent's set of goals.

We call this model Cooperative Individual Situation Assessment (CISA) between a group of agents. We define CISA as the exchange of information between friendly agents in order for each agent to individually perform unique situation assessment within its own "mind." This differs from other information sharing scenarios that could arise in a multiple agent setting by placing a certain constraints on the types of interactions agents can expect when sharing information. First, all agents can assume that the information received from other agents is qualified with respect to the task at hand rather than consisting of raw data. Secondly, all agents can assume friendly, rather than adversarial, interactions with other agents. By friendly interactions, all agents can assume that their communications with each other will follow the Cooperative Principle described by Grice [5] in the following set of supermaxims: (1) the

quantity of information provided is neither more or less informative than needed; (2) the **quality** of information does not include anything that the sending agent believes to be false or lacks adequate evidence; (3) the information exchanged is **relevant**; and (4) the **manner** in which the information is given avoids obscurity, ambiguity, and is brief and orderly.

This model makes no assumption of all the agents seeking a collective goal although it does allow for agents to do so if each agent individually holds the goal and decides to work toward it. For related material see [1,2,3,4,6,7,8].

2.0 The SAM-ISS Testbed

We have developed a prototype to simulate Surface to Air Missile (SAM) units comprised of three specialized vehicles. Within our testbed, the SAM-ISS prototype, we have implemented three modules: a domain simulator, a ISS performance analyzer, and a graphical user interface (GUI).

The testbed is currently modeled around an information sharing problem within the domain of Theater Missile Defense. This problem consists of three stationary sensing nodes on a battlefield that are each on the look out for Surface to Air Missile units (SAMs) comprised of three specific ground vehicles: a power vehicle, a communications vehicle, and a transporter vehicle.

Each sensing node is able to see only a portion of the battlefield, in a conical cross-section with the apex at the node, extending outward across the field, away from the node. This is marked on the battlefield as three triangles of visibility extending out from each of the three nodes. Due to the spacing of the nodes, some areas of the battlefield are out of view of all nodes, some areas are covered by one node, and other areas have redundant coverage. Band-shaped regions on the battlefield corresponding to the amount of node visibility coverage have been marked as low, medium, and high visibility. The high visibility region has complete and redundant coverage throughout the region, with over half of the region covered by 2 nodes. The medium visibility region has incomplete coverage with small pockets of redundant coverage and small pockets of no coverage. The low visibility region has large areas totally just over half of the region, with no visibility coverage; the other areas in the low visibility region have coverage only by a single node (see Figure 1).




Figure 1: Battlefield display showing nodes at the bottom of the field and the range of the node's sensors extending up from each node. Three visibility coverage regions have been marked on the battlefield based on the percentage of node sensor coverage within that area. These horizontal coverage regions have been marked as high, medium, and low visibility. Two groups of three vehicles, marked C1, T1, P1, and C2, T2, P2 are also shown. These represent the vehicles within two separate SAM units.

The six vehicles move through the field in accordance to one of several predefined scenarios. In each scenario, the vehicles travel in two groups of three, where each group is comprised of a transporter-erector-launcher vehicle (marked T), a communications vehicle (marked C), and a power vehicle (marked P). Such a grouping is called a SAM, a surface to air missile unit. These two SAMs can be positioned anywhere in the field prior to running a simulated scenario. As the vehicles move through the field during a scenario, the nodes detect the vehicles that exist within their field of view. If a node detects three vehicles moving in close formation such that the unit is comprised of a T, C, and P vehicle, the node will surmise that it is viewing a SAM unit.

A SAM unit can be in one of three modes while the simulation is running. These modes are moving, arming, and launching. When a SAM is arming or launching, that SAM is said to have had an event. Nodes that have detected a SAM are able to detect SAM events. Within this model we assume 100% sensor accuracy. If a vehicle is in a visibility region for a node, it will always be seen. Likewise, if the three vehicles comprising a SAM are detected together by one node within its visibility region, then the node realized that it has spotted a SAM.

3.0 Test Scenarios

Three scenarios for vehicle movement have been defined within this testbed, two of which have been used for the initial ISS analysis described within this paper. The first scenario is one in which one of the SAM units travels horizontally from left to right across the battlefield while the other unit travels horizontally from right to left. For this scenario, the SAM units can be placed within the low, medium, or high visibility bands, or in any combination of two bands. Within this scenario, a SAM unit stays within a single visibility band throughout the course of the scenario. The second scenario places one SAM unit such that it travels vertically through all

three visibility bands, from low to high while the other SAM unit travels horizontally across the battlefield within one of any of the three visibility bands.

These two scenarios allow for ISS tests in circumstances where this is little information to be shared up through scenarios where information could potentially be shared continuously.

In our implemented system we used the following specification as a specialized production rule:

```
IF <sender-filter>
THEN <sender> GENERATE
      FOR <recipient> THAT MEET <filter-potential-recipient>
SEND <filter-recipient>
```

The <sender-filter> is the set of conditions under which the message source will send a message. These conditions filter messages that the sender may send out. An example of specific condition is *the range data for TEL2 is less than 2 miles and gathered less than 2 seconds ago*. A more abstract example is *target must be identified*. In general, these are the (a) temporal conditions/constraints for communication between agents, (b) non-temporal conditions/constraints for data to be communicated between agents, e.g., reliability-threshold, information-age, availability. <sender> is the source of information, i.e., who sends the information. <message> is the composition format of the message, e.g., id:sender. <recipient> is the receiver of the message. It can be one or more sites. <filter-potential-recipient> are the conditions for selecting recipients known by all agents (i.e., common knowledge) on the recipients, such as the nodes that needs the info most, lacks the info the most, respond to solicitation, relevance-to-agent, e.g., *not busy or interested*. <filter-recipient> are conditions that the recipient might impose such as *no messages on TEL after 10pm*. The conditions in the rule encoded ISS teleologies.

Three information-sharing strategies have been defined among the nodes. These have been named: broadcast, event-driven, and unique. When using broadcast ISS, each node, on each time step, broadcasts all information it currently detects to all other nodes. Such information includes both observation of a type of vehicle and overall observation of a SAM unit. During event-driven ISS, a node will only broadcast information to all other nodes when a SAM is observed in the state of arming and launching. Both broadcast and event-driven ISS send information to all other nodes. Unique ISS differs from these by sending information only to predefined neighboring nodes. The teleological policies used by these three information sharing strategies are threefold. All three strategies share the same set of information: identification of each vehicle it is able to see, whether or not it has detected any SAMs, and whether or not any detected SAMs have been observed to be arming/launching. This information is always shared in an unsolicited manner such that an agent will share information without requiring a request from other agents. What differs between these three ISSs are who and when the information is shared. For both broadcast and event-based, the information is shared with everyone. For unique ISS, it is only shared with the predefined neighbors, limiting the number of messages sent and received during each phase of communication. Broadcast and unique ISS always share information any time anything of interest is noticed. Thus, each time an agent observes the battlefield, if a vehicle is identified, a SAM is detected, and/or an arming and

launching event is observed, this information is shared. Event-based differs by only sharing information when SAMs are observed having an arming and launching event, limiting the communications only to times when information of highest importance needs to be communicated.

Due to limited space we will omit details of simulator implementation. Our current implementation only includes a basic set of five ISS metrics that are gathered throughout the course of one running scenario and then displayed to the user. These metrics are:

- Timeliness - the amount of time before a node becomes aware of a certain piece of information;
- Consistency - the number of nodes aware of a certain piece of information;
- Completeness - the number of pieces of information that a node is aware of as a function of time;
- Cumulative consistency - the percentage of nodes aware of individual pieces of information as a function of time;
- Cumulative completeness - the average number of piece of information each node was aware of over time.

4.0 Results and Discussion

Data was collected for the first and second scenario using broadcast, event, and unique ISS such that all combinations of SAM travel within regions of visibility coverage where tested. The data collected was based on the five metrics discussed in the previous section: information timeliness, completeness, consistency, plus cumulative completeness and consistency. We omit details of experiments and scenarios due to space and only discuss results. Broadcast ISS, due to its policies of always sharing information with everyone, is able to maintain completeness and consistency of information across all nodes with the most information timeliness possible but at the cost of extremely high communications traffic. In fact, many of the messages sent using broadcast ISS are redundant as a node will send out information during every clock cycle that it detects anything with its sensors. Thus, if a node sees a communications vehicle during the first, second, and third cycles, it will send out messages to the other nodes on all three cycles saying that it saw a communications vehicle, despite the redundancy. In our current model, the agents do not possess any memory so this redundancy is necessary but in a scenario with more complex agents that have the ability to remember information, this ISS method would tend to provide an overwhelming amount of redundant information. This method also puts a greater strain on the communications network. The amount of information communicated throughout the course of a scenario is a function of:

$$(\text{cycles}) * (\text{no. nodes}) * (\text{no. nodes} - 1) * (\text{info packet size}) / (\% \text{ of visibility})$$

which is clearly dominated by the number of nodes participating in the scenario, giving a communications growth of N^2 where N equals the number of nodes.

Event-based ISS, which dictates that information is only shared when an event is observed, shows the performance of an ISS that is limited only to sharing information when something of critical importance occurs. This limits the amount of communications traffic considerably but at a cost to information timeliness, completeness, and consistency. If critical

events are difficult to detect because they are occurring in an area with very low visibility coverage, the level of information completeness and consistency drops considerably. In the case of events occurring at every time cycle, event-based ISS will behave exactly as broadcast ISS. Yet, the assumption of event-based ISS is that critical events are occurring only over a fraction of the total time. Information completeness and consistency will depend on the percentage of time that critical events occur and can be observed.

Unique ISS manages to reduce the amount of communications traffic between agents from broadcast ISS while providing higher information timeliness, completeness, and consistency than found with event-based ISS. The amount of information communicated over the course of a scenario is reduced only by sending to a maximum of two neighboring nodes. This reduces the communications growth to a linear function of the number of nodes. Information completeness and consistency will fair as well over large numbers of nodes but if the domain constraints can be satisfied by allowing neighborhoods of nodes to maintain local completeness and consistency, then this ISS methods performs better than the other two with respect to maintaining lower communications network traffic.

Of the set of metrics used in analysis, cumulative completeness, measuring the number of objects that each node was aware of over time, cumulative consistency, and measuring the percentage of time each node was aware of each SAM unit served to be accurate indicators of an ISSs ability to distribute information across all of the nodes. Additional metrics for collecting the amount of overall network traffic would be helpful in determining the cost of communications. It is possible that this could be used in conjunction with cumulative completeness and cumulative consistency to create a system that was able to choose an appropriate ISS on the fly based on the current domain constraints.

When vehicles traveled in low visibility areas, timeliness scores were consistently poorer. This information can be useful in meeting domain constraints. In such cases, a strategy such as event-based ISS would not perform well and a system could use such knowledge to switch to another ISS that performs better for completeness and consistency when vehicles are harder to detect.

5.0 Conclusion

We presented a model and a testbed for situation assessment. Many parameters affecting information sharing and strategies for information sharing were discussed in the form of generic questions. Clearly we are moving in the direction of developing a general framework that is applicable to a large class of problems. Our testbed is missing a few key parameters. If communication and deliberation did not cost time and effort, communication should be kept at maximum. However, in the real world, time and effort are critical factor. We showed a tradeoff of 15% detection rate for a 30% saving in communication. Cost and other factors such as models of other agents need to be included in the testbed to make it more useful.

Our approach assumes the cooperative principle and Gricean maxims used in both teaching and analyzing human discourse. However, in order for our information sharing strategies to be modeled such, we must add a great deal of reasoning and resource-sensitivity to

our agents. For example, we know what it means to be "relevant" but can not be modeled easily. Without attempting to solve the holy grail of AI we plan to (a) develop a more robust model of cooperative individual situation assessment, and (b) model simple forms of principles from language planning literature in linguistics.

6.0 References

1. MS Cohen, et al (1994) A Cognitive Framework for Battlefield Commander's Situation Assessment, **US Army Institute for Behavioral and Social Science**, TR 1002.
2. M.R. Endsley (1988). Design and Evolution for Situation Awareness enhancement. In 23rd annual **Human Factors Society**, Pages 97-101. Human Factors Society.
3. K. Decker (1995) Environment Centered Analysis and Design of Coordination Mechanisms, PhD Thesis, University of Massachusetts.
4. L. Foner (1995), Clustering and Information Sharing in an Ecology of Cooperative Agents, AAAI-95 Spring Symposium on **Information Gathering from Distributed Heterogeneous Environments**, Stanford University.
5. H.P. Grice (1975) Logic and Conversation, In P. Cole (editor) **Syntax and Semantics Vol 3**, pp. 41-58, Academic Press, New York.
6. Y. Lashkari, M. Metral, and P. Maes (1994) Collaborative Interface Agents, Proceedings of **AAAI-94 Conference**, Seattle, Washington 1994.
7. T. Malone, K. Grant, F. Turbak, S. Brodst, M. Cohen (1987). Intelligent Information-Sharing Systems, In *Computing Practices*, pages 390-402, ACM.
8. A. Moukas (1997) Amalthaea: Information Discovery and Filtering using Multiagent Evolving Ecosystem, **International Journal of Applied Artificial Intelligence** (to appear).
9. C. Smith, A. Sage (1991). A Theory of Situation Assessment for decision support. In *Information and Decision Technologies*, No. 17, pages 91-124. Elsevier Science Publishing, North Holland.

**MISSION
OF
AFRL/INFORMATION DIRECTORATE (IF)**

The advancement and application of Information Systems Science and Technology to meet Air Force unique requirements for Information Dominance and its transition to aerospace systems to meet Air Force needs.